

Simplicial Flat Norm with Scale

Sharif Ibrahim, Bala Krishnamoorthy*, Kevin R. Vixie

Department of Mathematics, Washington State University, Pullman, WA 99164.

{sibrahim,bkrishna,kvixie}@math.wsu.edu

Abstract

We study the multiscale simplicial flat norm (MSFN) problem, which computes flat norm at various scales of sets defined as oriented subcomplexes of finite simplicial complexes in arbitrary dimensions. We show that MSFN is NP-complete when homology is defined over integers. We cast MSFN as an instance of integer linear optimization. Following recent results on related problems, the MSFN integer program can be solved in polynomial time by solving its linear programming relaxation, when the simplicial complex satisfies a simple topological condition (absence of relative torsion). Our most significant contribution is the *simplicial deformation theorem*, which states that one may approximate a general current with a simplicial current while bounding the expansion of its mass. We present explicit bounds on the quality of this approximation, which indicate that the simplicial current gets closer to the original current as we make the simplicial complex finer. MSFN opens up the possibilities of using flat norm to denoise or extract scale information of large data sets in arbitrary dimensions. On the other hand, it allows one to employ the large body of algorithmic results on simplicial complexes to address more general problems related to currents.

1 Introduction

Currents are standard objects studied in geometric measure theory, and are named so by analogy with electrical currents that have a kind of magnitude and direction at every point. Intuitively, one could think of currents as generalized surfaces with orientations and multiplicities. The mathematical machinery of currents has been used to tackle many fundamental questions in geometric analysis, such as the ones related to area minimizing surfaces, isoperimetric problems, and soap-bubble conjectures [15]. We restrict our attention to currents with integer multiplicities.

The *mass* $M(T)$ of a current d -dimensional current T can be thought of intuitively as the weighted d -dimensional volume of the generalized object represented by T . For instance, the mass of a 2-dimensional current can be taken as the area of the surface it represents. The *flat norm* of a d -dimensional current T is given by

$$\mathbb{F}(T) = \min_S \{M(T - \partial S) + M(S) \mid T - \partial S \in \mathcal{E}_d, S \in \mathcal{E}_{d+1}\}, \quad (1)$$

where \mathcal{E}_d is the set of d -dimensional currents with compact support and integer multiplicities. One also uses flat norm to measure the “distance” between two d -currents. More precisely, the flat norm distance between two d -currents T and P is given by

$$\mathbb{F}(T, P) = \inf \{M(Q) + M(R) \mid T - P = Q + \partial R, Q \in \mathcal{E}_d, R \in \mathcal{E}_{d+1}\}. \quad (2)$$

*Corresponding author

Morgan and Vixie [16] showed that the L^1 total variation functional ($L^1\text{TV}$) introduced by Chan and Esedoğlu [3] computes the flat norm for boundaries T with integer multiplicity. Given this correspondence, and the use of *scale* in $L^1\text{TV}$, Morgan and Vixie defined [16] the flat norm with scale $\lambda \in [0, \infty)$ of an oriented d -dimensional set T as

$$\mathbb{F}_\lambda(T) \equiv \min_S \{V_d(T - \partial S) + \lambda V_{d+1}(S)\}, \quad (3)$$

where S varies over oriented $(d + 1)$ -dimensional sets, and V_d is the d -dimensional volume, used in place of mass. Figure 1 illustrates this definition. Flat norm of the 1D current T is given by the sum of the length of the resulting oriented curve $T - \partial S$ (shown separated from the input curve for clarity) and the area of the 2D patch S shown in red. Large values of λ , above the curvature of both humps in the curve T , preserve both humps. Values of λ between the two curvatures eliminate the hump on the right. Even smaller values “smooth out” both humps as illustrated here, giving a more “flat” curve, as S can now be comprised of much bigger 2D patches.




Figure 1: 1D current T , and flat norm decomposition $T - \partial S$ at appropriate scale λ . The resulting current is shown slightly separated from the input current for clearer visualization.

Figure 1 illustrates the utility of flat norm for deblurring or smoothing applications, e.g., in 3D terrain maps or 3D image denoising. But efficient methods for computing flat norm are known only for certain types of currents in two dimensions. For $d = 1$, Under the setting where T is a boundary, i.e., a loop, embedded in \mathbb{R}^2 and the minimizing surface $S \in \mathbb{R}^2$ as well, the flat norm could be calculated efficiently, for instance, using graph cut methods [13] – see the work of Goldfarb and Yin [11] and Vixie et al. [25], and references therein. Motivated by applications in image analysis, these approaches usually worked with a grid representation of the underlying space (\mathbb{R}^2). Pixels in the image readily provide such a representation.

While it is computationally convenient that $L^1\text{TV}$ minimizers give us the scaled flat norm for the input images, this approach restricts us to currents that are boundaries of codimension 1. Correspondingly, the calculation of flat norm for 1-boundaries embedded in higher dimensional spaces, e.g., \mathbb{R}^3 , or for input curves that are not necessarily boundaries has not received much attention so far. Similarly, flat norm calculations for higher dimensional input sets have also not been well-studied. Such situations often appear in practice – for instance, consider the case of an input set T that is a curve sitting on a manifold embedded in \mathbb{R}^3 , with choices for S restricted to this manifold as well. Further, computational complexity of calculating flat norm in arbitrary dimensions has not been studied. But this is not a surprising observation, given the continuous, rather than combinatorial, setting in which flat norm computation has been posed so far.

Simplicial complexes that triangulate the input space are often used as representations of manifolds. Such representations use triangular or tetrahedral meshes [9] as opposed to the uniform square or cubical grid meshes in \mathbb{R}^2 and \mathbb{R}^3 . Various simplicial complexes are often used to represent data (in any dimension) that captures interactions in a broad sense, e.g., the Vietoris–Rips complex to capture coverage of coordinate-free sensor networks [5, 6]. It is natural to consider flat norm calculations in such settings of simplicial complexes for denoising or regularizing sets, or for other similar tasks. At the same time, requiring that the simplicial complex be embedded in high dimensional space modeled by regular square grids may be cumbersome, and computationally prohibitive in many cases.

1.1 Our Contributions

We define a *simplicial* flat norm (SFN) for an input set T given as a subcomplex of the finite oriented simplicial complex K triangulating the set, or underlying space Ω . More generally, T is the simplicial representation of a rectifiable current with integer multiplicities. The choices of the higher dimensional sets S are restricted to K as well. We extend this definition to the multiscale simplicial flat norm (MSFN) by including a scale parameter λ . SFN is thus a special case of MSFN with the default value of $\lambda = 1$.

This discrete setting lets us address the worst case complexity of computing flat norm. Given its combinatorial nature, one would expect the problem to be difficult in arbitrary dimensions. Indeed, we show the problem of computing MSFN is NP-complete by reducing the optimal bounding chain problem (OBCP), which was recently shown to be NP-complete [8], to a special case of the MSFN problem. We cast the problem of finding the optimal S , and thus calculating MSFN, as an integer linear programming (ILP) problem. Given that the original problem is NP-complete, instances of this ILP could be hard to solve. Utilizing recent work [7] on the related optimal homologous chain problem (OHCP), we provide conditions on K under which this ILP problem can in fact be solved in polynomial time. In particular, MSFN can be computed in polynomial time when T is d -dimensional, and K is $(d + 1)$ -dimensional and orientable, for all $d \geq 0$. A similar result holds for the case when T is d -dimensional, and K is $(d + 1)$ -dimensional and embedded in \mathbb{R}^{d+1} , for all $d \geq 0$.

Our most significant contribution is the *simplicial deformation theorem* (Theorem 5.1), which states that given an arbitrary d -current in $|K|$ (underlying space), we are assured of an approximating current in the d -skeleton of K . This result is a substantial modification and generalization of the classical deformation theorem for currents on to square grids. Our deformation theorem explicitly specifies the dependence of the bounds of approximation on the regularity and size of the simplices in the simplicial complex. Hence it is immediate from the theorem that as we refine the simplicial complex K while preserving the bounds on simplicial regularity, the flat norm distance between an arbitrary d -current in $|K|$ and its deformation onto the d -skeleton of K vanishes.

1.2 Work on Related Problems

The problem of computing MSFN is closely related to two other problems on chains defined on simplicial complexes – the optimal homologous chain problem (OHCP) and the optimal bounding chain problem (OBCP). Given a d -chain \mathbf{t} of the simplicial complex K , the OHCP is to find a d -chain \mathbf{x} that is homologous to \mathbf{t} such that $\|\mathbf{x}\|_1$ is minimal. In the OBCP, we are given a d -chain \mathbf{t} of K , and the goal is to find a $(d + 1)$ -chain \mathbf{s} of K whose boundary is \mathbf{t} and $\|\mathbf{s}\|_1$ is minimal. The OBCP is closely related to the problem of finding an area-minimizing surface with a given boundary [15]. Computing the MSFN could be viewed, in a simple sense, as combining the objectives of the corresponding OHCP and OBCP instances, with the scale factor determining the relative importance of one objective over the other.

When \mathbf{t} is a cycle and the homology is defined over \mathbb{Z}_2 , Chen and Freedman showed that OHCP is NP-hard [4]. Dey, Hirani, and Krishnamoorthy [7] studied the original version of OHCP with homology defined over \mathbb{Z} , and showed that the problem is in fact solvable in polynomial time when K satisfies certain conditions (when it has no relative torsion). Recently, Dunfield and Hirani [8] have shown that the OHCP with homology defined over \mathbb{Z} is NP-complete. We will use their results to show that the problem of computing MSFN is NP-complete (see Section 2.1). These authors also showed that the OBCP with homology defined over \mathbb{Z} is NP-complete as well. Their result builds on the previous work of Agol, Hass, and Thurston [1], who showed that the knot genus problem is NP-complete, and a slightly different version of the least area surface problem is NP-hard.

The standard simplicial approximation theorem from algebraic topology describes how continuous maps are approximated by simplicial maps that satisfy the star condition [17, §14]. Our simplicial deformation

theorem applies to currents, which are more general objects than continuous maps. More importantly, we present explicit bounds on the expansion of mass of the current resulting from simplicial approximation. In his PhD thesis, Sullivan [22] considered deforming currents on to the boundary of convex sets in a cell complex, which are more general than the simplices we work with. But simplicial complexes admit efficient algorithms more naturally than cell complexes. We adopt a different approach for deformation than Sullivan. More significantly, our bounds on approximation are much smaller than the ones presented by Sullivan (see Section 5.2). Along with MSFN, our deformation theorem also establishes how the OHCP and OBCP could be used on general continuous inputs by taking simplicial approximations, thus expanding widely the applicability of this family of techniques to real-life data.

2 Definition of Simplicial Flat Norm

Consider a finite p -dimensional simplicial complex K triangulating the set Ω , where the simplices are oriented, with $p \geq d + 1$. The set T is defined as the integer multiple of an oriented d -dimensional subcomplex of K , representing a rectifiable d -current with integer multiplicity. Let m and n be the number of d - and $(d + 1)$ -dimensional simplices in K , respectively. The set T is then represented by the d -chain $\sum_{i=1}^m t_i \sigma_i$, where σ_i are all d -simplices in K and t_i are the corresponding *weights*. We will represent this chain by the vector of weights $\mathbf{t} \in \mathbb{Z}^m$. We use bold lower case letters to denote vectors, and the corresponding letter with subscript to denote components of the vector, e.g., $\mathbf{x} = [x_j]$. For \mathbf{t} representing the set T with integer multiplicity of one, $t_i \in \{-1, 0, 1\}$ with -1 indicating that the orientations of σ_i and T are opposite. But t_i can take any integer value in general. Thus, \mathbf{t} is the representation of T in the elementary d -chain basis of K . We consider $(d + 1)$ -chains in K modeling sets S representing rectifiable $(d + 1)$ -currents with integer multiplicities, and denote them similarly by $\sum_{j=1}^n s_j \tau_j$ in the elementary $(d + 1)$ -chain basis of K consisting of the individual simplices τ_j . We denote the chain modeling such a set S using the corresponding vector of weights $\mathbf{s} \in \mathbb{Z}^n$.

Relationships between the d - and $(d + 1)$ -chains of K are captured by its $(d + 1)$ -boundary matrix $[\partial_{d+1}]$, which is an $m \times n$ matrix with entries in $\{-1, 0, 1\}$. If the d -simplex σ_i is a face of the $(d + 1)$ -simplex τ_j , then the (i, j) entry of $[\partial_{d+1}]$ is nonzero, otherwise it is zero. This nonzero value is $+1$ if the orientations of σ_i and τ_j agree, and is -1 when their orientations are opposite. The d -chain representing the set $T - \partial_{d+1}S$ is then given as

$$\mathbf{x} = \mathbf{t} - [\partial_{d+1}]\mathbf{s}.$$

Notice that $\mathbf{x} \in \mathbb{Z}^m$. We define the simplicial flat norm (SFN) of T represented by the d -chain \mathbf{t} in the $(d + 1)$ -dimensional simplicial complex K as

$$F_S(T) = \min_{\mathbf{s} \in \mathbb{Z}^n} \left\{ \sum_{i=1}^m V_d(\sigma_i) |x_i| + \sum_{j=1}^n V_{d+1}(\tau_j) |s_j| \mid \mathbf{x} = \mathbf{t} - [\partial_{d+1}]\mathbf{s}, \mathbf{x} \in \mathbb{Z}^m \right\}. \quad (4)$$

Since \mathbf{x} and \mathbf{s} are chains in a simplicial complex, the masses of the currents they represent (as given in Equation 1) are indeed given by the weighted sums of the volumes of the corresponding simplices. The integer restrictions $\mathbf{x} \in \mathbb{Z}^m$ and $\mathbf{s} \in \mathbb{Z}^n$ are important in this definition. It is unclear how one would interpret a fractional weight for a simplex in these chains, as we are studying integral currents. We generalize the definition of SFN to define a *multiscale* simplicial flat norm (MSFN) of T in the simplicial complex K by including a scale parameter $\lambda \in [0, \infty)$.

$$F_S^\lambda(T) = \min_{\mathbf{s} \in \mathbb{Z}^n} \left\{ \sum_{i=1}^m V_d(\sigma_i) |x_i| + \lambda \left(\sum_{j=1}^n V_{d+1}(\tau_j) |s_j| \right) \mid \mathbf{x} = \mathbf{t} - [\partial_{d+1}]\mathbf{s}, \mathbf{x} \in \mathbb{Z}^m \right\}. \quad (5)$$

This definition is the simplicial version of the multiscale flat norm defined in Equation (3). The default, or nonscale, SFN in Equation (4) is a special case of MSFN with the default value of $\lambda = 1$. Henceforth we study the more general MSFN. We assume the d - and $(d + 1)$ -dimensional volumes of simplices to be any nonnegative values. For example, when σ_i is a 1-simplex, i.e., edge, $V_1(\sigma_i)$ could be taken as its Euclidean length. Similarly, $V_2(\tau_j)$ for a triangle τ_j could be its area. For ease of notation, we denote $V_d(\sigma_i)$ by w_i and $V_{d+1}(\tau_j)$ by v_j , with the dimensions d and $d + 1$ evident from the context.

Remark 2.1. The minimum in the definition of MSFN (Equation 5) indeed exists. The function

$$f^\lambda(T, S) = \sum_{i=1}^m w_i |x_i| + \lambda \left(\sum_{j=1}^n v_j |s_j| \right) \quad \text{with} \quad \mathbf{x} = \mathbf{t} - [\partial_{d+1}] \mathbf{s} \quad (6)$$

is lower bounded by zero, as it is the sum of nonnegative entries (we have $\lambda \geq 0$). Notice that $F_S^\lambda(T) = \min_S f^\lambda(T, S)$. Further, we only consider integral \mathbf{s} defined on the finite simplicial complex K , and hence there are only a finite number of values for this function. Hence its minimum indeed exists, which defines the MSFN of \mathbf{t} . On the other hand, the proof of existence of minimum in the original definition of flat norm for rectifiable currents employs the Hahn–Banach theorem [10, pg. 367].

We illustrate the optimal decompositions to compute the MSFN for two different scales ($\lambda = 1$ and $\lambda \ll 1$) in Figure 2. Notice that the input set T , shown in blue, is not a closed loop here. It is a subcomplex of the simplicial complex triangulating Ω . The underlying set Ω need not be embedded in \mathbb{R}^2 – it could be sitting in \mathbb{R}^3 or any higher dimension. We do not show the orientations of individual simplices and chains so as not to clutter the figure. We could take each triangle to be oriented counterclockwise (CCW), with T oriented CCW as well, and each edge oriented arbitrarily. When scale $\lambda = 1$, we get the default SFN of T , where the S chosen (shown in light pink) is such that the resulting optimal $T - \partial S$ (indicated by the thin curve in dark green) is devoid of all the “kinks”, but is similar to T in overall form. For $\lambda \ll 1$, the second term in the definition (Equation 5) contributes much less to the MSFN. As such, the optimal $T - \partial S$ consists of a short chain of two edges (shown in light green), which closes the original T curve to form a loop. S in this case includes the triangles in the former choice of S , and all other triangles enclosed by the original curve T and the resulting $T - \partial S$.




Figure 2: MSFN illustrated for two different scales $\lambda = 1$ and $\lambda \ll 1$. See text for explanation.

2.1 Complexity of MSFN

To study the complexity of computing the MSFN, we consider a decision version of the problem, termed *decision-MSFN* or DMSFN. The function $f^\lambda(T, S)$ used here is defined in Equation 6, with the modification that w_i and v_j are assumed to be rational for purposes of analyses of complexity.

Definition 2.2 (DMSFN). Given a p -dimensional finite simplicial complex K with $p \geq d + 1$, a set T defined as a d -subcomplex of K , a scale $\lambda \in [0, \infty)$, and a rational number $f_0 \geq 0$, does there exist a $(d + 1)$ -dimensional subcomplex S of K such that $f^\lambda(T, S) \leq f_0$?

The related optimal homologous chain problem (OHCP) was recently shown to be NP-complete [8, Theorem 1.4]. We reduce OHCP to a special case of DMSFN, thus showing that DMSFN is NP-complete as well. The default optimization version of MSFN consequently turns out to be NP-hard.

Theorem 2.3. DMSFN is NP-complete, and MSFN is NP-hard.

Proof. DMSFN lies in NP, as we can calculate $f^\lambda(T, S)$ in polynomial time when given a pair of d - and $(d + 1)$ -subcomplexes T and S , respectively, of the simplicial complex K . We just need to check that the chains \mathbf{t} and \mathbf{s} representing T and S satisfy $\mathbf{x} = \mathbf{t} - [\partial_{d+1}]\mathbf{s}$, which can be done in time polynomial in the dimensions m, n of $[\partial_{d+1}]$. Evaluating $f^\lambda(T, S)$ can also be done in time polynomial in m and n .

Conversely, we can reduce an equivalent instance of OHCP to a special case of DMSFN in a straightforward manner. The input d -chain for the OHCP instance is \mathbf{t} in the given simplicial complex K itself. Notice that the equations defining homologous chains in OHCP are given as $\mathbf{c} = \mathbf{t} + [\partial_{d+1}]\mathbf{s}$, which are equivalent to the decomposition equations in DMSFN, as reversing the sign of \mathbf{s} still gives a chain homologous to \mathbf{t} . Consider the special case of DMSFN with $\lambda = 0$ and $w_i = 1$ for $1 \leq i \leq m$. A solution to OHCP exists with the 1-norm of the homologous chain $\leq f_0$ if and only if this DMSFN instance has a solution $f^0(T, S) \leq f_0$. Since OHCP is NP-hard, we get that DMSFN is NP-hard as well. Therefore, DMSFN is NP-complete.

A polynomial time algorithm for the optimization version of MSFN problem implies one for DMSFN. Given an instance of DMSFN, we apply this algorithm to the corresponding optimization version, and compare the flat norm value to f_0 . Hence the MSFN problem is NP-hard. \square

We now consider attacking the MSFN problem using techniques from the area of discrete optimization. Even though the problem is NP-hard, this approach helps us to identify special cases in which we can compute the MSFN in polynomial time.

3 MSFN and Integer Linear Programming

The problem of finding the MSFN of the d -chain \mathbf{t} (Equation 5) can be cast formally as the following optimization problem.

$$\begin{aligned} & \text{minimize} && \sum_{i=1}^m w_i |x_i| + \lambda (\sum_{j=1}^n v_j |s_j|) \\ & \text{subject to} && \mathbf{x} = \mathbf{t} - [\partial_{d+1}]\mathbf{s}, \\ & && \mathbf{x} \in \mathbb{Z}^m, \mathbf{s} \in \mathbb{Z}^n. \end{aligned} \tag{7}$$

The objective function is piecewise linear in the integer variables \mathbf{x} and \mathbf{s} . Using standard modeling techniques from linear optimization [2, pg. 18], we can reformulate the problem as the following integer *linear* program (ILP).

$$\begin{aligned} & \min && \sum_{i=1}^m w_i (x_i^+ + x_i^-) + \lambda (\sum_{j=1}^n v_j (s_j^+ + s_j^-)) \\ & \text{s.t.} && \mathbf{x}^+ - \mathbf{x}^- = \mathbf{t} - [\partial_{d+1}](\mathbf{s}^+ - \mathbf{s}^-) \\ & && \mathbf{x}^+, \mathbf{x}^- \geq \mathbf{0}, \mathbf{s}^+, \mathbf{s}^- \geq \mathbf{0} \\ & && \mathbf{x}^+, \mathbf{x}^- \in \mathbb{Z}^m, \mathbf{s}^+, \mathbf{s}^- \in \mathbb{Z}^n. \end{aligned} \tag{8}$$

The objective function coefficients need to be nonnegative for this formulation to work – indeed, we have w_i, v_j , and λ nonnegative. Integer linear programming is NP-complete [18]. The linear programming relaxation of the ILP above is obtained by ignoring the integer restrictions on the variables.

$$\begin{aligned} \min \quad & \sum_{i=1}^m w_i(x_i^+ + x_i^-) + \lambda \left(\sum_{j=1}^n v_j(s_j^+ + s_j^-) \right) \\ \text{s.t.} \quad & \mathbf{x}^+ - \mathbf{x}^- = \mathbf{t} - [\partial_{d+1}](\mathbf{s}^+ - \mathbf{s}^-) \\ & \mathbf{x}^+, \mathbf{x}^- \geq \mathbf{0}, \quad \mathbf{s}^+, \mathbf{s}^- \geq \mathbf{0} \end{aligned} \tag{9}$$

We are interested in instances of this linear program (LP) that have integer optimal solutions, which hence are optimal solutions for the original ILP (Equation 8) as well. Totally unimodular matrices yield a prime class of linear programming problems with integral solutions. Recall that a matrix is totally unimodular if all its subdeterminants equal $-1, 0$, or 1 ; in particular, each entry is $-1, 0$, or 1 . The connection between total unimodularity and linear programming is specified by the following theorem.

Theorem 3.1. [24] *Let A be an $m \times n$ totally unimodular matrix, and $\mathbf{b} \in \mathbb{Z}^m$. Then the polyhedron $\mathcal{P} = \{\mathbf{x} \in \mathbb{R}^n \mid A\mathbf{x} = \mathbf{b}, \mathbf{x} \geq \mathbf{0}\}$ has integral vertices.*

Notice that the feasible set of the MSFN LP (Equation 9) has the form specified in the theorem above, with the variable vector $(\mathbf{x}^+, \mathbf{x}^-, \mathbf{s}^+, \mathbf{s}^-)$ in place of \mathbf{x} . The corresponding equality constraint matrix A has the form $[I \quad -I \quad B \quad -B]$, where I is the identity matrix and $B = [\partial_{d+1}]$. The input d -chain \mathbf{t} is in place of the right-hand side vector \mathbf{b} . In order to use Theorem 3.1 for computing MSFN, we connect the total unimodularity of constraint matrix A and that of boundary matrix B .

Lemma 3.2. *If $B = [\partial_{d+1}]$ is totally unimodular, then so is the matrix $A = [I \quad -I \quad B \quad -B]$.*

Proof. Starting with B , we get the matrix A by appending columns of B scaled by -1 to its right, and appending columns with a single nonzero entry of ± 1 to its left. Both these classes of operations preserve total unimodularity [18, pg. 280]. \square

Consequently, we get the following result on polynomial time computability of MSFN.

Theorem 3.3. *If the boundary matrix $[\partial_{d+1}]$ of the finite oriented simplicial complex K is totally unimodular, then the multiscale simplicial flat norm of the set T specified as a d -chain $\mathbf{t} \in \mathbb{Z}^m$ of K can be computed in polynomial time.*

Proof. The problem of computing the MSFN of T (Equation 5) is cast as the optimization problem given in Equation (7). This problem is reformulated as an instance of ILP (Equation 8). We get the MSFN LP (Equation 9) by relaxing the integrality constraints of this ILP. As noted in Remark 2.1, the optimal cost of this LP is finite. The polyhedron of this LP has at least one vertex, given that all variables are nonnegative [2, Cor. 2.2]. By Lemma 3.2, the constraint matrix of this LP is totally unimodular, as $[\partial_{d+1}]$ is so. Hence by Theorem 3.1, all vertices of the feasible region of MSFN LP are integral, since $\mathbf{t} \in \mathbb{Z}^m$.

An optimal solution $(\mathbf{x}_*^+, \mathbf{x}_*^-, \mathbf{s}_*^+, \mathbf{s}_*^-)$ of the MSFN LP can be found in polynomial time using an interior point method [2, Chap. 9]. If it happens to be a unique optimal solution, then it will be a vertex, and hence will be integral by Theorem 3.1. Hence it is an optimal solution to the ILP (Equation 8).

If the optimal solution is not unique, then $(\mathbf{x}_*^+, \mathbf{x}_*^-, \mathbf{s}_*^+, \mathbf{s}_*^-)$ may be nonintegral. But since the optimal cost is finite, there must exist a vertex in its polyhedron that has this minimum cost. Given a nonintegral optimal solution obtained by an interior point method, one can find such an integral optimal solution at a vertex in polynomial time [12]. Hence the MSFN ILP can be solved in polynomial time in this case as well. \square

Remark 3.4. We point out that since the boundary matrix $B = [\partial_{d+1}]$ has entries only in $\{-1, 0, 1\}$, the constraint matrix of the MSFN LP (Equation 9) also has entries only in $\{-1, 0, 1\}$. Hence the MSFN LP can be solved in strongly polynomial time [23], i.e., the time complexity is independent of the objective function and right-hand side coefficients, and depends only on the dimensions of the problem.

Remark 3.5. Components of variables $\mathbf{x}^+, \mathbf{x}^-, \mathbf{s}^+, \mathbf{s}^-$ in the MSFN ILP (Equation 8) could assume values other than $\{-1, 0, 1\}$, indicating integer multiplicities higher than 1 for the corresponding simplices in the optimal decomposition. The definition of MSFN (Equation 5) does allow such larger multiplicities. At the same time, if one insists on using each $(d+1)$ -simplex at most *once* when calculating MSFN, and insists on similar restrictions on d -simplices in the optimal decomposition, we can modify the ILP such that Theorem 3.3 still holds.

Denoting the entire variable vector by $\mathbf{x} = (\mathbf{x}^+, \mathbf{x}^-, \mathbf{s}^+, \mathbf{s}^-) \in \mathbb{Z}^{2m+2n}$, we add the upper bound constraints $\mathbf{x} \leq \mathbf{1}$, where $\mathbf{1}$ is the $(2m+2n)$ -vector of ones. These inequalities could be converted to the set of equations $\mathbf{x} + \mathbf{y} = \mathbf{1}$, where \mathbf{y} is the $(2m+2n)$ -vector of slack variables that are nonnegative. These modifications give an ILP whose polyhedron is in the same form as described in Theorem 3.1, with the equations denoted as $A'\mathbf{x}' = \mathbf{b}'$ for the variable vector $\mathbf{x}' = (\mathbf{x}, \mathbf{y})$. The new constraint matrix A' is related to the constraint matrix A of the original MSFN ILP given in Lemma 3.2 as

$$A' = \begin{bmatrix} A & O \\ I & I \end{bmatrix},$$

where I is the $2m+2n$ identity matrix, and O is the $m \times (2m+2n)$ zero matrix. Hence A' is obtained from A by first adding $2m+2n$ rows with a single nonzero entry of $+1$, and then adding to the resulting matrix $2n+2m$ more columns with a single nonzero entry of $+1$. These operations preserve total unimodularity [18, pg. 280], and hence the new constraint matrix A' is totally unimodular when $[\partial_{d+1}]$ is so. The new right-hand side vector $\mathbf{b}' \in \mathbb{Z}^{3m+2n}$ consists of the input chain \mathbf{t} and the vector of ones from the new upper bound constraints.

Since the efficient computability of MSFN depends on the total unimodularity of the boundary matrix, we study the conditions under which total unimodularity of boundary matrices can be guaranteed.

4 Simplicial Complexes and Relative Torsion

Dey, Hirani, and Krishnamoorthy [7] have given a simple characterization of the simplicial complex whose boundary matrix is totally unimodular. In short, if the simplicial complex does not have relative torsion then its boundary matrix is totally unimodular. We state this and other related results here for the sake of completeness, and refer the reader to the original paper [7] for details and proofs. The simplicial complex K in these results has dimension $d+1$ or higher. Recall that a d -dimensional simplicial complex is *pure* if it consists of d -simplices and their faces, i.e., there are no lower dimensional simplices that are not part of some d -simplex in the complex.

Theorem 4.1. [7, Theorem 5.2] *The boundary matrix $[\partial_{d+1}]$ of a finite simplicial complex K is totally unimodular if and only if $H_d(L, L_0)$ is torsion-free for all pure subcomplexes L_0, L of K , with $L_0 \subset L$.*

These authors further describe situations in which the absence of relative torsion is guaranteed. The following two special cases describe simplicial complexes for which the boundary matrix is always totally unimodular.

Theorem 4.2. [7, Theorem 4.1] *The boundary matrix $[\partial_{d+1}]$ of a finite simplicial complex triangulating a compact orientable $(d+1)$ -dimensional manifold is totally unimodular.*

Theorem 4.3. [7, Theorem 5.7] *The boundary matrix $[\partial_{d+1}]$ of a finite simplicial complex embedded in \mathbb{R}^{d+1} is totally unimodular.*

For simplicial complexes of dimension 2 or lower, the boundary matrix is totally unimodular when the complex does not have a Möbius subcomplex.

Theorem 4.4. [7, Theorem 5.13] *For $d \leq 1$, the boundary matrix $[\partial_{d+1}]$ is totally unimodular if and only if the finite simplicial complex has no $(d + 1)$ -dimensional Möbius subcomplex.*

We illustrate the implications of the results above for the efficient computation of MSFN by considering certain sets. When the input set T is of dimension 1, and is described on an orientable 2-manifold to which the choices of 2-dimensional set S are also restricted, we can always compute its MSFN by solving the MSFN LP (Equation 9) in polynomial time. A similar result holds when T is a set of dimension 2 described as a subcomplex of a 3-complex sitting in \mathbb{R}^3 . For a 1-dimensional set T with choices of S restricted to a 2-complex K , we can always compute the MSFN of T efficiently as long as K does not have a 2-dimensional Möbius subcomplex. Notice that K itself need not be embedded in \mathbb{R}^3 for this result to work – it could be sitting in some higher dimensional space.

5 Simplicial Deformation Theorem

The deformation theorem [10, Sections 4.2.7–9] is one of the fundamental results of geometric measure theory, and more particularly of the theory of currents. It approximates an integral current by deforming it onto a cubical grid of appropriate mesh size. On the other hand, we have been studying currents or sets in the setting of simplicial complexes, rather than on square grids. Our proof is a substantial modification of the classical proof of the deformation theorem. We found the presentation of the latter proof by Krantz and Parks [14, Section 7.7] especially helpful. Our proof mimics their proof when possible. The gist of this theorem is the assertion that we may approximate a current with a simplicial current.

Recall that $V_d(\sigma)$ denotes the d -dimensional volume of a d -simplex σ . The *perimeter* of σ is the set of all its $(d - 1)$ -dimensional faces, denoted as $\text{Per}(\sigma) = \{\cup_j \tau_j \mid \tau_j \in \sigma, \dim(\tau_j) = d - 1\}$. We will also refer to the $(d - 1)$ -dimensional volume of $\text{Per}(\sigma)$ as the perimeter of σ , but denote it as $P(\sigma) = \sum_{\tau_j \in \text{Per}(\sigma)} V_{d-1}(\tau_j)$. We let $D(\sigma)$ be the diameter of σ , which is the largest Euclidean distance between any two points in σ .

Theorem 5.1 (Simplicial Deformation Theorem). *Let K be a p -dimensional simplicial complex embedded in \mathbb{R}^q , with $p = d + k$ for $k \geq 1$ and $q \geq p$. Suppose that for every simplex $\sigma \in K$*

$$\frac{D(\sigma) P(\sigma)}{V_d(\mathcal{B}_\sigma)} \leq \kappa_1 < \infty,$$

$$\frac{D(\sigma)}{r_\sigma} \leq \kappa_2 < \infty,$$

and

$$D(\sigma) \leq \Delta$$

hold, where $\hat{\mathcal{B}}_\sigma$ is the largest ball inscribed in σ , \mathcal{B}_σ is the ball with half the radius and same center as $\hat{\mathcal{B}}_\sigma$, and r_σ is the radius of \mathcal{B}_σ . Let T be a d -dimensional current in \mathbb{R}^q such that $T \subset |K|$, i.e., T is a subset of the underlying space of K . Suppose that T satisfies

$$M(T) + M(\partial T) < \infty.$$

Then there exists a simplicial d -current P supported in the d -skeleton of K whose boundary ∂P is supported in the $(d - 1)$ -skeleton of K such that

$$T - P = Q + \partial R,$$

and the following controls on mass M hold:

$$M(P) \leq (4\vartheta_K)^k M(T) + \Delta(4\vartheta_K)^{k+1} M(\partial T), \quad (10)$$

$$M(\partial P) \leq (4\vartheta_K)^{k+1} M(\partial T), \quad (11)$$

$$M(R) \leq \Delta(4\vartheta_K)^k M(T), \text{ and} \quad (12)$$

$$M(Q) \leq \Delta(4\vartheta_K)^k (1 + 4\vartheta_K) M(\partial T), \quad (13)$$

where $\vartheta_K = \kappa_1 + \kappa_2$.

Remark 5.2. It is immediate that the flat norm distance between T and P can be made arbitrarily small by subdividing the simplicial complex to reduce Δ while preserving the regularity of the refinement as measured by κ_1 and κ_2 .

Remark 5.3. Note that this theorem combines the unscaled and scaled versions of the original deformation theorem [14, Theorems 7.7.1 and 7.7.2] into one theorem through the explicit form of the constraints. In our proof of Theorem 5.1, we replace certain pieces of the original proof as presented by Krantz and Parks [14, Pages 211–222] without reproducing all the other details of their proof. We found their exposition quite well-structured, making it easier to identify the modifications needed to get our theorem.

Remark 5.4. The bound for $M(P)$ in Theorem 5.1 is larger than the classical bound. We get this large bound because we generate P through retractions alone, and not using the usual Sobolev-type estimates [14, Pages 220–222]. And of course, the Δ in the coefficient of the extra term means that it becomes unimportant as the simplicial complex is subdivided.

5.1 Proof of the Simplicial Deformation Theorem

At the heart of the modification of the deformation theorem (from cubical grid to simplicial complex settings) is the recalculation of an integral over the current and its boundary. This integral appears in a bound on the *Jacobian of the retraction*, which measures the expansion in mass of the current resulting from the process of retracting it on to the simplices of the simplicial complex. To do this recalculation, we consider the retraction ϕ one step at a time, building it through independent choices of centers to project from in every simplex and its every face.

We first describe the general set up of retraction within a simplex. We then present certain bounds on the mass expansion resulting from the retraction in Lemmas 5.6, 5.7, and 5.8. In particular, we obtain bounds on the expansion that are independent of the choice of points from which we project. These bounds are independent of the particular current that we retract on to the simplicial complex. But we employ these bounds to subsequently bound the overall expansion of mass of the current resulting from the retraction.

5.1.1 Retracting from a center inside a simplex

We describe the details of retraction for an ℓ -simplex σ in the p -dimensional simplicial complex K . This set up is valid for any ℓ , but in particular, we will use the bounds thus obtained for $d \leq \ell \leq p$ when retracting a d -current onto K . We pick a center $\mathbf{a} \in \text{Int}(\sigma)$, the interior of σ , and project every $\mathbf{x} \in \text{Int}(\sigma) \setminus \{\mathbf{a}\}$ along the ray $(\mathbf{x} - \mathbf{a})/\|\mathbf{x} - \mathbf{a}\|$ to $\text{Per}(\sigma)$. Denoting this map as $\phi(\mathbf{x}, \mathbf{a})$, we get

$$\phi(\mathbf{x}, \mathbf{a}) = (\phi_\pi \circ \phi_\delta)(\mathbf{x}, \mathbf{a}), \quad (14)$$

where $\phi_\delta(\mathbf{x}, \mathbf{a})$ is a dilation of \mathbb{R}^ℓ by the factor $\|\phi(\mathbf{x}, \mathbf{a}) - \mathbf{a}\|/\|\mathbf{x} - \mathbf{a}\|$ and $\phi_\pi(\mathbf{x}, \mathbf{a})$ is a nonorthogonal projection along $(\mathbf{x} - \mathbf{a})/\|\mathbf{x} - \mathbf{a}\|$ onto $\tau_{\mathbf{x}}$, the $(\ell - 1)$ -dimensional face of σ containing $\phi(\mathbf{x}, \mathbf{a})$. We denote $\hat{r} = \|\phi(\mathbf{x}, \mathbf{a}) - \mathbf{a}\|$ and $r = \|\mathbf{x} - \mathbf{a}\|$. Let E_ℓ be the ℓ -hyperplane that contains σ and $E_{\ell-1}$ the $(\ell - 1)$ -hyperplane that contains $\tau_{\mathbf{x}}$. Denote the orthogonal projection of \mathbf{a} onto $E_{\ell-1}$ by \mathbf{b} , and let $\hat{h} = \|\mathbf{b} - \mathbf{a}\|$. For any point $\mathbf{y} = \mathbf{a} + (\mathbf{b} - \mathbf{a})\gamma$ with $0 < \gamma < 1$, we get $\phi(\mathbf{y}, \mathbf{a}) = \mathbf{b}$. In particular, we consider the point of intersection of line connecting \mathbf{a} and \mathbf{b} with the $(\ell - 1)$ -hyperplane parallel to $\tau_{\mathbf{x}}$ that contains \mathbf{x} . Naming this point \mathbf{y} , we define $h = \|\mathbf{y} - \mathbf{a}\|$. Let $\mathbf{z} \in E_\ell$ denote either normal to $\tau_{\mathbf{x}}$ at $\phi(\mathbf{x}, \mathbf{a})$ (either of the two possibilities work). Let $\mathbf{v}_2 = (\mathbf{x} - \mathbf{a})/\|\mathbf{x} - \mathbf{a}\|$, and let \mathbf{v}_1 be the vector in $\text{span}(\mathbf{z}, \mathbf{v}_2)$ that is normal to \mathbf{v}_2 and points into σ . We illustrate this construction on a 3-simplex in Figure 3, where the cone of \mathbf{a} with face τ is shown in red and the other points and vectors are labeled. We also illustrate the corresponding slice spanned by \mathbf{v}_1 and \mathbf{v}_2 in Figure 4.




Figure 3: Illustration of the dilation and nonorthogonal projection involved in retraction for a 3-simplex.




Figure 4: A 2-dimensional illustration of the dilation calculation.

Choose an orthogonal basis $\{\mathbf{w}_1, \dots, \mathbf{w}_{\ell-2}\}$ for $\{\text{span}(\mathbf{v}_1, \mathbf{v}_2)\}^\perp$. Note that $\text{span}(\mathbf{w}_1, \dots, \mathbf{w}_{\ell-2}) \in$

$E_{\ell-1}$. Let \mathbf{w}' be a unit vector in $\{\text{span}(\mathbf{w}_1, \dots, \mathbf{w}_{\ell-2})\}^\perp \cap E_{\ell-1}$ parallel to $\phi(\mathbf{x}) - \mathbf{b}$. Then $\{\mathbf{v}_1, \mathbf{w}_1, \dots, \mathbf{w}_{\ell-2}, \mathbf{v}_2\}$ is an orthogonal basis for \mathbb{R}^ℓ , and ϕ_π is given by

$$\begin{aligned}\phi_\pi(\mathbf{v}_1) &= \alpha \mathbf{w}', \\ \phi_\pi(\mathbf{w}_i) &= \mathbf{w}_i, \quad i \in \{1, \dots, \ell-2\}, \quad \text{and} \\ \phi_\pi(\mathbf{v}_2) &= 0,\end{aligned}\tag{15}$$

where $\alpha = \hat{r}/\hat{h}$. Notice that the above set up works everywhere except when $\phi(\mathbf{x}) = \mathbf{b}$, in which case we obtain an orthogonal projection for $\phi_\pi(\mathbf{x})$ along $\mathbf{b} - \mathbf{a}$. Choosing coordinates for the tangent spaces of σ and $\tau_{\mathbf{x}}$ to be $\{\mathbf{v}_1, \mathbf{w}_1, \dots, \mathbf{w}_{\ell-2}, \mathbf{v}_2\}$ and $\{\mathbf{w}', \mathbf{w}_1, \dots, \mathbf{w}_{\ell-2}\}$, respectively, we get from Equation (15) that $D\phi_\pi(\mathbf{x}, \mathbf{a})$ is the $(\ell-1) \times \ell$ matrix given as

$$D\phi_\pi(\mathbf{x}, \mathbf{a}) = \begin{bmatrix} \alpha & 0 & 0 & \dots & 0 & 0 \\ 0 & 1 & 0 & \dots & 0 & 0 \\ \vdots & & & \ddots & & \\ 0 & \dots & & & 1 & 0 \end{bmatrix}.\tag{16}$$

5.1.2 Bounding the Integral of the Jacobian

We now present a series of bounds on integrals of the dilation of d -volumes induced by the retraction. Since $\ell = d$ implies we are already in the d -skeleton and no retraction is needed, we can assume that $\ell > d$. We start with a bound on the maximum dilation of d -volumes under the retraction ϕ . $D\phi$ will denote the *tangent map* or *Jacobian map* of ϕ .

Definition 5.5. Let $J_d\phi(\mathbf{x}, \mathbf{a})$ be the maximum dilation of d -volumes induced by $D\phi(\mathbf{x}, \mathbf{a})$ at \mathbf{x} .

We will use the definitions and results on $D\phi_\pi(\mathbf{x}, \mathbf{a})$ in ℓ -dimension given above. In particular, recall what $D(\sigma)$, \hat{h} , and h stand for.

Lemma 5.6. For any center \mathbf{a} and any point $\mathbf{x} \neq \mathbf{a}$ in the ℓ -simplex σ with $d < \ell \leq p = d + k$,

$$J_d\phi(\mathbf{x}, \mathbf{a}) \leq \left(\frac{\hat{h}}{h}\right)^d \frac{D(\sigma)}{\hat{h}}.$$

Proof. Following Equation (14), we seek bounds on $D\phi_\delta(\mathbf{x}, \mathbf{a})$ and $D\phi_\pi(\mathbf{x}, \mathbf{a})$. Since $D\phi_\delta(\mathbf{x}, \mathbf{a})$ simply scales by $\hat{r}/r = \hat{h}/h$, the expansion of d -volume of any d -hyperplane by $D\phi_\delta(\mathbf{x}, \mathbf{a})$ is by a factor of $(\hat{h}/h)^d$. On the other hand, bounding the dilation that $D\phi_\pi(\mathbf{x}, \mathbf{a})$ can cause in d -hyperplanes is a little more involved. We seek a bound on

$$\frac{\sqrt{\det((D\phi_\pi(\mathbf{x}, \mathbf{a})U)^T(D\phi_\pi(\mathbf{x}, \mathbf{a})U))}}{\sqrt{\det(U^TU)}}\tag{17}$$

for all $\ell \times d$ matrices U . Using the generalized Pythagorean theorem [14, Section 1.5], we get

$$\det(U^TU) = \sum_{\lambda \in \Lambda} (\det(U_\lambda))^2$$

where submatrix U_λ consists of the d rows of U specified by the set of index maps Λ given as

$$\lambda \in \Lambda \equiv \{f|f : [1, \dots, d] \rightarrow [1, \dots, \ell], \text{ } f \text{ is one to one and increasing}\}.$$

A similar result holds for $\det((D\phi_\pi(\mathbf{x}, \mathbf{a})U)^T(D\phi_\pi(\mathbf{x}, \mathbf{a})U))$, with the functions f considered mapping $[1, \dots, d]$ to $[1, \dots, \ell-1]$.

Observe that multiplying by $D\phi_\pi(\mathbf{x}, \mathbf{a})$ (Equation 16) just scales the first row of U by α and removes the last row. Thus $\alpha \det(U_\lambda) \geq \det((D\phi_\pi(\mathbf{x}, \mathbf{a})U)_\lambda)$, which implies that α is a bound on the ratio in Equation (17). Thus we have that

$$J_d\phi(\mathbf{x}, \mathbf{a}) = \left(\frac{\hat{r}}{r}\right)^d \frac{\|\phi(\mathbf{x}, \mathbf{a}) - \mathbf{a}\|}{\|\mathbf{b} - \mathbf{a}\|} = \left(\frac{\hat{h}}{h}\right)^d \frac{\|\phi(\mathbf{x}, \mathbf{a}) - \mathbf{a}\|}{\|\mathbf{b} - \mathbf{a}\|} \leq \left(\frac{\hat{h}}{h}\right)^d \frac{D(\sigma)}{\hat{h}}$$

holds for all \mathbf{x} and \mathbf{a} in σ , where $D(\sigma)$ is the diameter of the ℓ -simplex σ . \square

Next we describe a bound on the integral of $J_d\phi(\mathbf{x}, \mathbf{a})$ over the entire ℓ -simplex, for a fixed center \mathbf{a} . We will find that this bound is independent of the position of \mathbf{a} . Recall that $\text{Per}(\sigma)$ and $\text{P}(\sigma)$ denote the perimeter of ℓ -simplex σ and the $(\ell - 1)$ -dimensional volume of the perimeter, respectively, and $\text{Int}(\sigma)$ its interior.

Lemma 5.7. *For any fixed center \mathbf{a} in the ℓ -simplex σ with $d < \ell \leq p = d + k$,*

$$\int_{\text{Int}(\sigma)} J_d\phi(\mathbf{x}, \mathbf{a}) d\mathcal{L}^\ell(\mathbf{x}) \leq D(\sigma) \text{P}(\sigma).$$

Proof. Consider the $(\ell - 1)$ -dimensional faces τ_j of σ , with $\text{Per}(\sigma) = \{\cup_j \tau_j \mid \tau_j \in \sigma, \dim(\tau_j) = \ell - 1\}$. Let σ_j denote the ℓ -simplex generated by \mathbf{a} and τ_j . Then

$$\int_{\text{Int}(\sigma)} J_d\phi(\mathbf{x}, \mathbf{a}) d\mathcal{L}^\ell(\mathbf{x}) = \sum_j \int_{\text{Int}(\sigma_j)} J_d\phi(\mathbf{x}, \mathbf{a}) d\mathcal{L}^\ell(\mathbf{x}).$$

Let $\tau_j(h)$ denote the $(\ell - 1)$ -simplex formed by the intersection of σ_j and the $(\ell - 1)$ -hyperplane parallel to τ_j at a distance h from \mathbf{a} . Thus, $\tau_j(\hat{h})$ is τ_j itself. We observe that our bound on $J_d\phi(\mathbf{x}, \mathbf{a})$ is constant in $\tau_j(h)$ for any h . The $(\ell - 1)$ -dimensional volume of $\tau_j(h)$ is given by

$$\text{V}_{\ell-1}(\tau_j(h)) = \left(\frac{h}{\hat{h}}\right)^{\ell-1} \text{V}_{\ell-1}(\tau_j).$$

Using the bound on $J_d\phi(\mathbf{x}, \mathbf{a})$ from Lemma 5.6, and noting that $D(\sigma_j) \leq D(\sigma) \forall j$, we get

$$\int_{\text{Int}(\sigma_j)} J_d\phi(\mathbf{x}, \mathbf{a}) d\mathcal{L}^\ell(\mathbf{x}) \leq \int_0^{\hat{h}} \left(\frac{h}{\hat{h}}\right)^{\ell-1} \text{V}_{\ell-1}(\tau_j) \left(\frac{\hat{h}}{h}\right)^d \frac{D(\sigma)}{\hat{h}} dh = \frac{\text{V}_{\ell-1}(\tau_j) D(\sigma)}{\ell - d}.$$

Summing this quantity over all $\tau_j \in \text{Per}(\sigma)$ and replacing $\ell - d \geq 1$ with 1 gives the overall bound. \square

We now bound the integral of $J_d\phi(\mathbf{x}, \mathbf{a})$ over centers \mathbf{a} with a fixed \mathbf{x} that we are retracting onto $\text{Per}(\sigma)$. Examination of the corresponding proof for the original deformation theorem [14, Section 7.7] shows that symmetry of the cubical mesh plays a very special role, which cannot be duplicated in the case of simplicial complex. In particular, we must avoid integrating over \mathbf{a} close to the perimeter of σ . Hence we integrate over as big a region as we can while still avoiding a neighborhood of the perimeter. As in the statement of the main Theorem 5.1, let $\hat{\mathcal{B}}_\sigma$ be the largest ball inscribed in σ , \mathcal{B}_σ be the ball with half the radius and same center as $\hat{\mathcal{B}}_\sigma$, and r_σ be the radius of \mathcal{B}_σ .

Lemma 5.8. *For any point \mathbf{x} in the ℓ -simplex σ with $d < \ell \leq p = d + k$,*

$$\int_{\mathcal{B}_\sigma} J_d\phi(\mathbf{x}, \mathbf{a}) d\mathcal{L}^\ell(\mathbf{a}) \leq D(\sigma) \text{P}(\sigma) + \text{V}_\ell(\mathcal{B}_\sigma) \frac{D(\sigma)}{r_\sigma}.$$

Proof. Similar to the subsimplices of σ considered in the Proof of Lemma 5.7, let σ_j now denote the ℓ -simplex formed by \mathbf{x} and $\tau_j \in \text{Per}(\sigma)$. In order to derive an upper bound, we integrate instead over regions that are by construction bigger than these subsimplices of σ . Denoting the simplex σ_j as Region 1, we define Regions 2 and 3 as follows. We refer the reader to Figure 5 for an illustration of this construction. Let σ'_j be the reflection of σ_j through \mathbf{x} , and similarly, let τ'_j be the reflection through \mathbf{x} of τ_j . We define σ'_j as Region 2. Notice that unlike Region 1, Region 2 need not be contained fully in σ . As defined in Section 5.1.1, let \mathbf{z} be the unit vector normal to the $(\ell - 1)$ -hyperplane containing τ_j pointing into σ . We define Region 3 as the ℓ -dimensional set $\tau'_j + [0, D(\sigma)]\mathbf{z}$, as illustrated in Figure 5.

Note that the union of all Region 2's and Region 3's cover σ . By an argument almost identical to that above, we have the following upper bound on the integrand in question.

$$\left(\frac{h'}{h}\right)^d \frac{\|\phi(\mathbf{x}, \mathbf{a}) - \mathbf{a}\|}{\hat{h}} \leq \left(\frac{h'}{h}\right)^d \frac{D(\sigma)}{\hat{h}}.$$

Integrating the second of these two terms over Region 2 and summing the integral over all such Regions




Figure 5: Illustration for integration of Jacobian bound over centers \mathbf{a} instead of \mathbf{x} 's. For the case of the triangle shown, there will be 3 sets of 3 regions. In general there will be 3 regions for every face of the simplex.

2 for all faces τ_j , we get the upper bound of $D(\sigma) P(\sigma)$. Here we use the same arguments as the ones employed in Lemma 5.7. Region 2 alone is not guaranteed to cover \mathcal{B}_σ as some of \mathcal{B}_σ may occupy parts of Region 3. Since $\mathbf{a} \in \mathcal{B}_\sigma$, we have $\hat{h} > r_\sigma$, and $h' \leq h$ when $\mathbf{a} \in \text{Region 3}$, so that

$$\left(\frac{h'}{h}\right)^d \frac{D(\sigma)}{\hat{h}} \leq \frac{D(\sigma)}{r_\sigma}.$$

Combining the above estimates while integrating over all such Regions 2 and 3 gives us the bound specified in the Lemma. \square

5.1.3 Bounding the pushforwards of the current

We consider the d -current T , and employ the bounds on the Jacobian of retraction described above to the pushforwards of T and its boundary ∂T on to the simplicial complex K . Our treatment of the pushforwards essentially follows the corresponding results of Krantz and Parks for the case of square grid [14, Pages 218–219]. We denote by $\|T\|$ the total variation measure of the current T .

Lemma 5.9. *Suppose K is a p -dimensional simplicial complex with $p = d + k$ for $k \geq 1$. Consider the stepwise retraction of the d -current $T \subset K$ (the $(d - 1)$ -current $\partial T \subset K$) onto the d -skeleton of K (respectively, the $(d - 1)$ -skeleton of K). Each step of the retraction on to the perimeter of an ℓ -simplex σ for $d < \ell \leq p$ (respectively, $d \leq \ell \leq p$) increases the mass of T or ∂T by at most a factor of*

$$4\vartheta_K = 4(\kappa_1 + \kappa_2) = 4 \max_{\sigma \in K} \left(\frac{D(\sigma) P(\sigma)}{V_\ell(\mathcal{B}_\sigma)} + \frac{D(\sigma)}{r_\sigma} \right).$$

Proof. Using Fubini's theorem [14, Page 26] and applying the bound in Lemma 5.8, we get

$$\int_{\mathcal{B}_\sigma} \int_\sigma J_d \phi(\mathbf{x}, \mathbf{a}) d\|T\|(\mathbf{x}) d\mathcal{L}^\ell(\mathbf{a}) = \int_\sigma \int_{\mathcal{B}_\sigma} J_d \phi(\mathbf{x}, \mathbf{a}) d\mathcal{L}^\ell(\mathbf{a}) d\|T\|(\mathbf{x}) \leq \vartheta_\sigma M(T|_\sigma),$$

where $\vartheta_\sigma = D(\sigma) P(\sigma) + V_d(\mathcal{B}_\sigma)(D(\sigma)/r_\sigma)$ and $T|_\sigma$ is the portion of the current T restricted to the simplex σ . Consider the subset of \mathcal{B}_σ defined as

$$H_T = \left\{ \mathbf{a} \in \mathcal{B}_\sigma \mid \int_\sigma J_d \phi(\mathbf{x}, \mathbf{a}) d\|T\|(\mathbf{x}) > \frac{4\vartheta_\sigma M(T|_\sigma)}{V_\ell(\mathcal{B}_\sigma)} \right\}.$$

Then $V_\ell(H_T) \leq (1/3) V_\ell(\mathcal{B}_\sigma)$. Similarly we define $H_{\partial T}$ for the pushforward of ∂T and get $V_\ell(H_{\partial T}) \leq (1/3) V_\ell(\mathcal{B}_\sigma)$. Then the set $\mathcal{B}_\sigma \setminus \{H_T \cup H_{\partial T}\}$ defines a subset of \mathcal{B}_σ with positive measure, with the centers \mathbf{a} in this subset satisfying $\int_\sigma J_d \phi(\mathbf{x}, \mathbf{a}) d\|T\|(\mathbf{x}) \leq 4\vartheta_\sigma M(T|_\sigma)/V_\ell(\mathcal{B}_\sigma)$ and $\int_\sigma J_d \phi(\mathbf{x}, \mathbf{a}) d\|\partial T\|(\mathbf{x}) \leq 4\vartheta_\sigma M(\partial T|_\sigma)/V_\ell(\mathcal{B}_\sigma)$. Hence we can choose centers to retract from in each simplex σ such that the expansion of mass of the current restricted to that simplex is bounded by $4\vartheta_\sigma/V_\ell(\mathcal{B}_\sigma)$. The bound specified in the Lemma follows when we consider retracting the entire current over multiple simplices in K , and set $\vartheta_K = \max_{\sigma \in K} \vartheta_\sigma$ as the generic upper bound that holds for all simplices in K . \square

Bound on complete sequence of retractions. We can apply the bound specified in Lemma 5.9 over multiple levels ℓ . Pushing T onto the d -skeleton of p -complex K multiplies the mass of T by a factor of at most $(4\vartheta_K)^k$. Likewise, pushing ∂T on to the $(d - 1)$ -skeleton multiplies the mass of ∂T by a factor of at most $(4\vartheta_K)^{k+1}$.

5.1.4 Bounding the distance between the current and its simplicial approximation

In the final step, we construct the simplicial current P approximating the original current T , and bound the flat norm distance between the two. Since we are now considering retraction maps over many simplices simultaneously, we refer to all such maps as ϕ in general, suppressing the particular \mathbf{x} and \mathbf{a} . Hence we map T forward in k applications of ϕ , picking centers (see Lemma 5.8) to project from in each step and in each simplex. We denote the composition of all these steps as ψ^1 . We pick each of these centers such that the retractions map ∂T with bounded amplification of mass as well (see Lemma 5.9).

Define the homotopy $g(\gamma, \mathbf{x}) = \gamma\mathbf{x} + (1 - \gamma)\psi^1(\mathbf{x})$ for $\gamma \in [0, 1]$. Then the homotopy formula [14, Section 7.4.3] gives

$$T = \psi_{\#}^1(T) + \partial g_{\#}([0, 1] \times T) + g_{\#}([0, 1] \times \partial T).$$

We define $R = g_{\#}([0, 1] \times T)$ and $Q_1 = g_{\#}([0, 1] \times \partial T)$. Then we get

$$T - \psi_{\#}^1(T) = \partial R + Q_1. \quad (18)$$

Finally, we map $\psi_{\#}^1(\partial T)$ forward to the $(d-1)$ -skeleton of simplicial complex K with one more application of ϕ to get $\psi_{\#}^2(\partial T) = \phi_{\#}(\psi_{\#}^1(\partial T))$. For this purpose, consider the homotopy $h(\gamma, \mathbf{x})$ from $\psi_{\#}^1(\partial T)$ to $\psi_{\#}^2(\partial T)$, i.e.,

$$h(\gamma, x) = \gamma\psi_{\#}^1(\mathbf{x}) + (1 - \gamma)\psi_{\#}^2(\mathbf{x}) \text{ for } \gamma \in [0, 1].$$

We define

$$P = \psi_{\#}^1(T) - h_{\#}([0, 1] \times \psi_{\#}^1(\partial T)). \quad (19)$$

P is a d -current whose boundary ∂P is contained in the $(d-1)$ -skeleton of K . Define $Q_2 = h_{\#}([0, 1] \times \psi_{\#}^1(\partial T))$. Using the homotopy formula, we get

$$\begin{aligned} \partial P &= \partial (\psi_{\#}^1(T) - h_{\#}([0, 1] \times \psi_{\#}^1(\partial T))) \\ &= \psi_{\#}^1(\partial T) - \partial h_{\#}([0, 1] \times \psi_{\#}^1(\partial T)) \\ &= \psi_{\#}^2(\partial T) \subset (d-1)\text{-skeleton of } K. \end{aligned}$$

Equation (19) gives $\psi_{\#}^1(T) = P + Q_2$. Defining $Q = Q_1 + Q_2$, Equation (18) gives

$$\begin{aligned} T - (P + Q_2) &= \partial R + Q_1, \text{ hence} \\ T - P &= \partial R + Q. \end{aligned}$$

Finally, we apply the bounds on the retraction described in Lemma 5.9 and the paragraph following this Lemma to the masses of the pushforwards. Noticing that $D(\sigma) \leq \Delta$ for all $\sigma \in K$, we get the following bounds, which finish the proof of our simplicial deformation theorem (Theorem 5.1).

$$\begin{aligned} M(P) &\leq (4\vartheta_K)^k M(T) + \Delta(4\vartheta_K)^{k+1} M(\partial T) \\ &= (4\vartheta_K)^k (M(T) + \Delta(4\vartheta_K) M(\partial T)), \\ M(\partial P) &\leq (4\vartheta_K)^{k+1} M(\partial T), \\ M(R) &\leq \Delta M(\psi_{\#}^1(T)) \\ &\leq \Delta(4\vartheta_K)^k M(T), \text{ and} \\ M(Q) &\leq \Delta(4\vartheta_K)^k M(\partial T) + \Delta(4\vartheta_K)^{k+1} M(\partial T) \\ &= \Delta(4\vartheta_K)^k (1 + 4\vartheta_K) M(\partial T). \end{aligned}$$

□

Remark 5.10. The influence of *Simplicial regularity* as measured by κ_1 and κ_2 is clearly revealed by the statement of our deformation theorem (Theorem 5.1). Explicit constants are a simple yet useful part of the result; as observed above in Remark 5.2, the statement of this theorem leads to an easy observation that the flat norm distance between T and P can be made as small as desired by subdividing the simplicial complex without increasing the bound on regularity.

Remark 5.11. We did not explicitly discuss the case of 0-dimensional currents. In this case, the bounds on mass expansion are all equal to one.

5.2 Comparison of Bounds of Approximation

Sullivan studied the deformation of integral currents on to the skeleton of a cell complex, which is composed of compact convex sets. He presented a deformation theorem for deforming integral currents on to the boundary of a cell complex [22, Theorem 4.5]. For ease of comparison, we use *our* notation to restate the bounds given by Sullivan for deforming a d -current T to a polyhedral current P in the boundary of a cell complex in \mathbb{R}^q . Recall that in our simplicial deformation theorem (Theorem 5.1), the simplicial complex considered has dimension p , and is embedded in \mathbb{R}^q for $q \geq p$.

$$M(P) \leq \binom{q}{d} \left(2d \left(\frac{d+1}{2d} \kappa_2 \right)^{d+1} \right)^{q-d+1} M(T), \quad (20)$$

$$M(\partial P) \leq \binom{q}{d-1} \left(2d \left(\frac{d+1}{2d} \kappa_2 \right)^d \right)^{q-d+1} M(\partial T), \quad \text{and} \quad (21)$$

$$\mathbb{F}(T, P) \leq (q - d + 1) \Delta (M(P) + M(\partial P)). \quad (22)$$

Our results corresponding to the first two bounds in Equations (20) and (21) are presented in Equations (10) and (11) in Theorem 5.1. We use $k = p - d$ to denote the codimension, and thus $k = p - d \leq q - d$. Hence the exponents of κ_2 terms in these bounds of Sullivan are at least $(k+1)d$, while those in our bounds are only k or $k+1$. To obtain the flat norm distance corresponding to the third bound given by Sullivan in Equation (22), we use the definition of flat norm distance between two currents specified in Equation (2). Using $T - P = \partial Q + R$, we combine two of our bounds specified in Equations (12) and (13) to get

$$\mathbb{F}(T, P) \leq \Delta (4\vartheta_K)^k (M(T) + (1 + 4\vartheta_K) M(\partial T)).$$

Once again, the exponents of regularity terms κ_1 and κ_2 in our bound are at most $k+1$, while those in Sullivan's bound are at least $(k+1)d$.

The differences in bounds highlighted above become stark when we are studying high-dimensional currents with high codimension. We get much smaller bounds because we deform in a different way. Also, we work with simplicial complexes while Sullivan uses the more general cell complexes. At the same time, efficient algorithms are more naturally constructed for simplicial complexes than for cell complexes.

6 Computational Results

We illustrate computations of MSFN by describing the flat norm decompositions of a 2-manifold with boundary embedded in \mathbb{R}^3 (see Figure 6). The input set has the underlying shape of a pyramid, to which several peaks and troughs of varying scale, as well as random noise, have been added. We model this set as a piecewise linear 2-manifold with boundary, and find a triangulation of the same as a subcomplex of a tetrahedralization of the $2 \times 2 \times 2$ cube centered at the origin, within which the set is located. We use the method of constrained Delaunay tetrahedralization [21] implemented in the package TetGen [19] for this purpose. We then compute the MSFN decomposition of the input set at various scale (λ) values. At high values, e.g., when $\lambda = 6$, the optimal decomposition resembles the input set with the small kinks due to random noise smoothed out. At the other end, for $\lambda = 0.01$, the optimal decomposition resembles a flat “sheet”. For intermediate values of λ , the optimal decomposition captures features of the input set at varying scales.




Figure 6: Top left: A view of original pyramidal surface in three dimensions. The remaining three figures show the flat norm decomposition for scales $\lambda = 6$ (top right), $\lambda = 2$ (bottom left), and $\lambda = 0.01$ (bottom right). See text for further explanation. The images were generated using the package TetView [20].

The entire 3-complex mesh modeling the cube in question consisted of 14,002 tetrahedra and 28,844 triangles. For each λ , computation of MSFN described above took only a few minutes on a regular PC using standard functions from MATLAB. This example demonstrates the feasibility of efficiently computing flat norm decompositions of large datasets in high dimensions, for the purposes of denoising or to recover scale information of the data.

7 Discussion

Our result on simplicial deformation (Theorem 5.1) places the definition of MSFN into clear context. If a current lives in the underlying space of a simplicial complex, we can deform it to be a simplicial current on the simplicial complex, and do so with *controlled* error. In fact, by subdividing the simplicial complex carefully, we can move this error as close to zero as we like. Since MSFN could be computed efficiently when the simplicial complex does not have relative torsion, one could naturally use our approach to compute the flat norm of a large majority of currents in arbitrarily large dimensions.

The MSFN problem, similar to the recent results on the OBCP [8], apply notions from algebraic topology and discrete optimization to problems from geometric measure theory such as flat norm of currents and area-minimizing hypersurfaces. What other classes of problems from the broader area of geometric analysis could we tackle using similar approaches? One such question appears to be the following. Is the flat norm decomposition of an integral current always another integral current? If not, could we specify conditions under which this result holds? Working in the setting of simplicial complexes, results on the existence of integral optimal solutions for instances of ILPs with integer right-hand side vectors may prove useful in answering this question.

While L^1 TV and flat norm computations have been used widely on data in two dimensions, such as images, the MSFN opens up the possibility of utilizing flat norm computations for higher dimensional data. Similar to the flat norm-based signatures for distinguishing shapes in two dimensions [25], could we define shape signatures using MSFN computations to characterize the geometry of sets in arbitrary dimensions? The sequence of optimal MSFN decompositions of a given set for varying values of the scale parameter λ captures all the scale information of its geometry. Could we represent all this information in a compact manner, for instance, in the form of a barcode?

Acknowledgments

Vixie and Ibrahim acknowledge financial support from National Science Foundation through the grant DMS-0914809.

References

- [1] Ian Agol, Joel Hass, and William Thurston. The computational complexity of knot genus and spanning area. *Transactions of the American Mathematical Society*, 358(9):3821–3850, 2006.
- [2] Dimitris Bertsimas and John N. Tsitsiklis. *Introduction to Linear Optimization*. Athena Scientific, Belmont, MA., 1997.
- [3] Tony F. Chan and Selim Esedoglu. Aspects of total variation regularized L^1 function approximation. *SIAM Journal on Applied Mathematics*, 65(5):1817–1837, 2005.
- [4] Chao Chen and Daniel Freedman. Hardness results for homology localization. In *SODA '10: Proc. 21st Ann. ACM-SIAM Sympos. Discrete Algorithms*, pages 1594–1604, 2010.

- [5] Vin de Silva and Robert Ghrist. Coordinate-free coverage in sensor networks with controlled boundaries via homology. *International Journal of Robotics Research*, 25(12):1205–1222, 2006.
- [6] Vin de Silva and Robert Ghrist. Coverage in sensor networks via persistent homology. *Algebraic and Geometric Topology*, 7:339–358, 2007.
- [7] Tamal K. Dey, Anil N. Hirani, and Bala Krishnamoorthy. Optimal homologous cycles, total unimodularity, and linear programming. In *STOC '10: Proc. 42nd Ann. Sympos. Theo. Comput.*, pages 221–230, 2010.
- [8] Nathan M. Dunfield and Anil N. Hirani. The least spanning area of a knot and the optimal bounding chain problem. In *ACM Symposium on Computational Geometry (SoCG 2011)*, pages 135–144, 2011.
- [9] Herbert Edelsbrunner. *Geometry and Topology for Mesh Generation*. Cambridge Monographs on Applied and Computational Mathematics. Cambridge University Press, New York, NY, USA, 2006.
- [10] Herbert Federer. *Geometric Measure Theory*. Die Grundlehren der mathematischen Wissenschaften, Band 153. Springer-Verlag, New York, 1969.
- [11] Donald Goldfarb and Wotao Yin. Parametric maximum flow algorithms for fast total variation minimization. *SIAM Journal on Scientific Computing*, 31(5):3712–3743, 2009.
- [12] Osman Güler, Dick den Hertog, Cornelis Roos, Tamas Terlaky, and Takashi Tsuchiya. Degeneracy in interior point methods for linear programming: a survey. *Annals of Operations Research*, 46-47(1):107–138, March 1993.
- [13] Vladimir Kolmogorov and Ramin Zabih. What energy functions can be minimized via graph cuts? In *Proceedings of the 7th European Conference on Computer Vision-Part III, ECCV '02*, pages 65–81, London, UK, UK, 2002. Springer-Verlag.
- [14] Steven G. Krantz and Harold R. Parks. *Geometric Integration Theory*. Cornerstones. Birkhauser, 2008.
- [15] Frank Morgan. *Geometric measure Theory: A Beginner's Guide*. Academic Press, fourth edition, 2008.
- [16] Simon P. Morgan and Kevin R. Vixie. L^1 TV computes the flat norm for boundaries. *Abstract and Applied Analysis*, 2007:Article ID 45153,14 pages, 2007.
- [17] James R. Munkres. *Elements of Algebraic Topology*. Addison–Wesley Publishing Company, Menlo Park, 1984.
- [18] Alexander Schrijver. *Theory of Linear and Integer Programming*. Wiley-Interscience Series in Discrete Mathematics. John Wiley & Sons Ltd., Chichester, 1986.
- [19] Hang Si. TetGen: A quality tetrahedral mesh generator and a 3d Delaunay triangulator. Available at <http://tetgen.berlios.de>.
- [20] Hang Si. TetView: A tetrahedral mesh and piecewise linear complex viewer. Available at <http://tetgen.berlios.de/tetview.html>.
- [21] Hang Si. Constrained delaunay tetrahedral mesh generation and refinement. *Finite Elements in Analysis and Design*, 46:33–46, January 2010.

- [22] John M. Sullivan. *A Crystalline Approximation Theorem for Hypersurfaces*. PhD thesis, Princeton University, 1990.
- [23] Éva Tardos. A strongly polynomial algorithm to solve combinatorial linear programs. *Operations Research*, 34(2):250–256, March 1986.
- [24] Arthur F. Veinott, Jr. and George B. Dantzig. Integral extreme points. *SIAM Review*, 10(3):371–372, 1968.
- [25] Kevin R. Vixie, Keith Clawson, Thomas J. Asaki, Gary Sandine, Simon P. Morgan, and Brandon Price. Multiscale flat norm signatures for shapes and images. *Applied Mathematical Sciences*, 4(13-16):667–680, 2010.